

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
15 November 2001 (15.11.2001)

PCT

(10) International Publication Number
WO 01/86454 A2(51) International Patent Classification⁷: G06F 13/00

DORMITZER, Paul, H.; c/o ADC Inc., 12501 Whitewater Drive, Minnetonka, MN 55343 (US).

(21) International Application Number: PCT/US01/14810

(74) Agents: CHOW, Stephen, Y. et al.; Perkins, Smith & Cohen, LLP, One Beacon Street, 30th floor, Boston, MA 02108 (US).

(22) International Filing Date: 8 May 2001 (08.05.2001)

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/566,540 8 May 2000 (08.05.2000) US

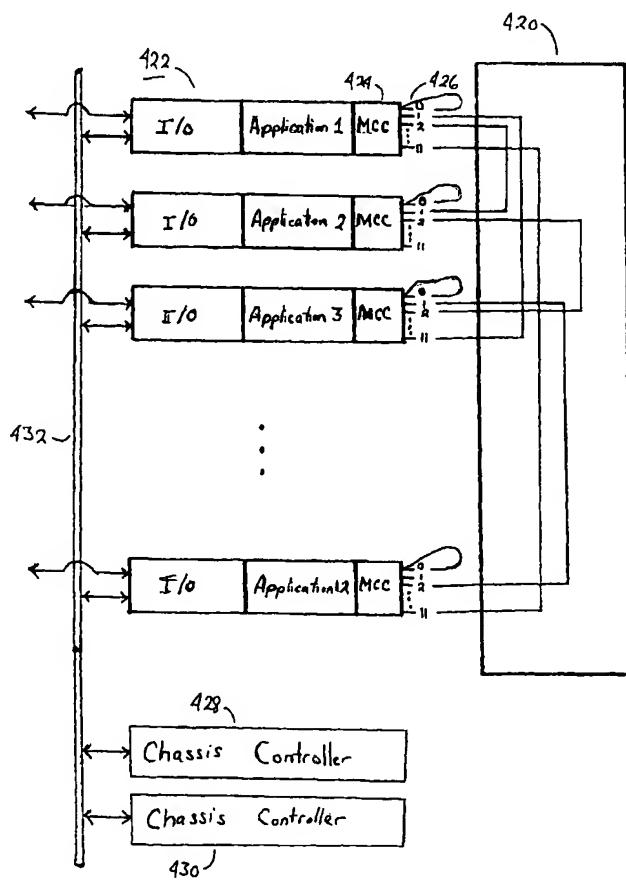
(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,

(71) Applicant: ADC INC. [US/US]; 12501 Whitewater Drive, Minnetonka, MN 55343 (US).

(72) Inventors: CHINNASWAMY, Kumar; c/o ADC Inc., 12501 Whitewater Drive, Minnetonka, MN 55343 (US).

[Continued on next page]

(54) Title: SYSTEM HAVING A MESSED BACKPLANE AND PROCESS FOR TRANSFERRING DATA THERE THROUGH



(57) Abstract: A meshed backplane has dedicated pairs of connections for high-speed serial connection in each direction between each of multiple application modules and each other application module. A management/control bus is provided for out-of-band signaling. The mesh of serial differential pairs may be used for management/control bus signals when necessary. A time division multiplexing fabric is also provided for telephony applications. A star interconnection region is provided for distribution of signals from redundant clocks.

WO 01/86454 A2



IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

- *without international search report and to be republished upon receipt of that report*

SYSTEM HAVING A MESHED BACKPLANE AND PROCESS FOR TRANSFERRING DATA THERETHROUGH

FIELD OF THE INVENTION

5 This invention relates generally to networking data processing systems, and, more particularly to a broadband network environment, for example, one using a SONET backbone and Hybrid Fiber-Coax(ial cable) ("HFC") to connect users to the backbone. An emerging hardware/software standard for the
10 HFC environment is DOCSIS (Data Over Cable Service Interface Standard) of CableLabs.

BACKGROUND OF THE INVENTION

15 In a current configuration of wide-area delivery of data to HFC systems (each connected to 200 households/clients), the head-end server is connected to a SONET ring via a multiplexer drop on the ring (see Fig. 1). These multiplexers currently cost some \$50,000 in addition to the head-end server, and scaling up of service of a community may require new
20 multiplexers and servers.

The failure of a component on the head-end server can take an entire "downstream" (from the head-end to the end-user) sub-network out of communication with the world.

25 Attempts have been made to integrate systems in order to reduce costs and to ease system management. A current integrated data delivery system is shown in Figure 2. Figure 2 shows a system having a reverse path monitoring system, an ethernet switch, a router, modulators and upconverters, a provisioning system, telephony parts, and a plurality of
30 CMTS's (cable modem termination systems). This type of system typically has multiple vendors for its multiple systems, has different management systems, a large footprint, high power requirements and high operating costs.

35 A typical network broadband cable network for delivery of voice and data is shown in Figure 3. Two OC-12 port interface servers are each connected to one of two backbone routers which are in turn networked to two switches. The switches are networked to CMTS head-end routers. The CMTS head-end routers

are connected to a plurality of optical nodes. The switches are also connected to a plurality of telephone trunk gateways which are in turn connected to the public switched telephone network (PSTN). As with the "integrated" system shown in 5 Figure 2, this type of network also typically has multiple vendors for its multiple systems, has different management systems, a large footprint, high power requirements and high operating costs.

In order to facilitate an effective integrated solution, 10 an integrated electrical interconnect system is needed that provides the bandwidth and routing capability for delivery of voice and data in a compact area.

Problems of integration include difficulties in circuit routing in a complex system, maintaining flexibility in 15 arrangement and type of components, managing data flow control, limiting component size and number and limiting cost.

It is desirable to have an integrated solution to reduce the size of the system, its power needs and its costs, as well as to give the data delivery system greater consistency.

20 It is an object of the present invention to provide a system and process for electrical interconnect for broadband delivery of high-quality voice, data, and video services.

It is another object of the present invention to provide 25 a system and process for a cable access platform having high network reliability with the ability to reliably support lifeline telephony services and the ability to supply tiered voice and data services.

It is another object of the present invention to provide a system and process for a secure, scalable network switch.

30

SUMMARY OF THE INVENTION

The problems of providing the bandwidth and routing capability for delivery of voice and data in a compact area 35 for an integrated switch are solved by the present invention of a meshed backplane.

The meshed backplane has dedicated pairs of connections for high-speed serial connection in each direction between

each of multiple application modules and each other application module. A management/control bus is provided for out-of-band signaling. The mesh of serial differential pairs may be used for management/control bus signals when necessary.

5 A time division multiplexing fabric is also provided for telephony applications. A star interconnect region is provided for distribution of signals from redundant clocks.

The present invention together with the above and other advantages may best be understood from the following detailed 10 description of the embodiments of the invention illustrated in the drawings, wherein:

BRIEF DESCRIPTION OF THE DRAWINGS

15 Fig. 1 shows a prior art network on a SONET ring;

Fig. 2 shows a prior art data delivery system;

Fig. 3 shows a prior art data delivery network;

Fig. 4 is a block diagram of a chassis according to principles of the invention;

20 Fig. 5 shows an integrated cable infrastructure having the chassis of Fig. 4;

Fig. 6 is a block diagram of the application cards, the backplane and a portion of the interconnections between them in the chassis of Fig. 4;

25 Fig. 7 is a schematic diagram of the backplane interconnections, including the switching mesh;

Fig. 8 is a block diagram of two exemplary slots showing differential pair connections between the slots;

30 Fig. 9 is a block diagram of the MCC chip in an application module according to principles of the present invention;

Fig. 10 is a diagram of a packet tag;

Fig. 11 is a block diagram of a generic switch packet header;

35 Fig. 12 is a flow chart of data transmission through the backplane;

Fig. 13 is a block diagram of an incoming ICL packet;

Fig. 14 is a block diagram of a header for the ICL packet of Fig. 13;

Fig. 15 shows example mapping tables mapping channels to backplane slots according to principles of the present
5 invention;

Fig. 16 is a block diagram of a bus arbitration application module connected in the backplane of the present invention; and

Fig. 17 is a state diagram of bus arbitration in the
10 application module of Fig. 16.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Figure 4 shows a chassis 200 operating according to principles of the present invention. The chassis 200 integrates a plurality of network applications into a single switch system. The invention is a fully-meshed OSI Layer 3/4 IP-switch with high performance packet forwarding, filtering and QoS/CoS (Quality of Service/Class of Service) capabilities using low-level embedded software controlled by a cluster
15 manager in a chassis controller. Higher-level software resides in the cluster manager, including router server functions (RIPv1, RIPv2, OSPF, etc.), network management (SNMP V1/V2), security, DHCP, LDAP, and remote access software (VPNs, PPTP, L2TP, and PPP), and can be readily modified or upgraded.

In the present embodiment of the invention, the chassis 200 has fourteen (14) slots for modules. Twelve of those fourteen slots hold application modules 205, and two slots hold chassis controller modules 210. Each application module has an on-board DC-DC converter and is "hot-pluggable" into
25 the chassis. The chassis controller modules 210 are for redundant system clock/bus arbitration. Examples of applications that may be integrated in the chassis are a CMTS module 215, an Ethernet module 220, a SONET module 225, and a telephony application 230. Another application may be an
30 interchassis link (ICL) port 235 through which the chassis may be linked to another chassis.

Fig. 5 shows an integrated cable infrastructure 260 having the chassis 200 of Fig. 4. The chassis 200 is part of

a regional hub 262 (also called the "head-end"), for voice and data delivery. The hub 262 includes a video controller application 264, a video server 266, Web/cache servers 268, and an operation support system (OSS) 270, a combiner 271 and 5 the chassis 200. The chassis 200 acts as an IP access switch. The chassis 200 is connected to a SONET ring 272, outside the hub 262, having a connection to the Internet 274, and a connection to the Public Switched Telephone Network (PSTN) 276. The chassis 200 and the video-controller application 264 10 are attached to the combiner 271. The combiner 271 is connected by an HFC link 278 to cable customers and provides IP voice, data, video and fax services. At least 2000 cable customers may be linked to the head-end by the HFC link 278. The chassis 200 can support a plurality of HFC links and also 15 a plurality of chassis may be networked together (as described below) to support thousands of cable customers.

By convention today, there is one wide-band channel for transmission (downloading) to users (which may be desktop computers, facsimile machines or telephone sets) and four much 20 narrower channels for (uploading). This is processed by the HFC cards with duplexing at an O/E node. The local HFC cable system or loop may be a coaxial cable distribution network with a drop to a cable modem.

Figure 6 shows application modules connected to a 25 backplane 420 of the chassis 200 of Figure 4. In the present embodiment of the invention, the backplane is implemented as a 24-layer printed wiring board and includes 144 pairs of uni-directional differential-pair connections, each pair directly connecting input and output terminals of each of a maximum of 30 twelve application modules with output and input terminals of each other module and itself. Each application module interfaces with the backplane through a Mesh Communication Chip (MCC) 424 through these terminals. Each application module is also connected to a chassis management bus 432 which 35 provides the modules with a connection to the chassis controllers 428, 430. Each MCC 424 has twelve (12) serial link interfaces that run to the backplane 420. Eleven of the serial links on each application module are for connecting the

application module to every other application module in the chassis. One link is for connecting the module with itself, i.e., a loop-back. The backplane is fully meshed meaning that every application module has a direct link to every other 5 application module in the chassis through the serial links. Only a portion of the connections is shown in Figure 6 as an example. The backplane mesh is shown in Figure 7.

The 12 channels with serial links of the MCC are numbered 0 to 11. This is referred to as the channel ID or CID. The 10 slots on the backplane are also numbered from 0 to 11 (slot ID, or SID). The chassis system does not require, however, that a channel 0 be wired to a slot 0 on the backplane. A serial link may be connected to any slot. The slot IDs are dynamically configured depending on system topology. This 15 provides freedom in backplane wiring layout which might otherwise require layers additional to the twenty-four layers in the present backplane. The application module reads the slot ID of the slot into which it is inserted. The application module sends that slot ID out its serial lines in 20 an idle stream in between data transmissions. The application module also includes the slot ID in each data transmission.

Figure 15 shows examples of mapping tables of channels in cards to backplane slots. Each card stores a portion of the table, that is, the table row concerning the particular card. 25 The table row is stored in the MCC.

Figure 16 shows a management bus arbitration application module connected to the backplane. The backplane contains two separate management buses for failure protection. Each application module in the chassis including the two chassis 30 controllers as well as the twelve applications modules can use both or either management bus.

The management bus is used for low-speed data transfer within the chassis and generally consists of control, statistical, configuration information, data from the chassis 35 controller modules to the application modules in the chassis.

The implementation of the management bus consists of a four bit data path, a transmit clock, a transmit clock signal, a collision control signal, and a four bit arbitration bus.

As seen in Figure 16, the bus controller has a 10/100 MAC device, a receive FIFO, bus transceiver logics, and a programmable logic device ("PLD").

The data path on the management bus is a four-bit (Media 5 Independent Interface) MII standard interface for 10/100 ethernet MACs. The bus mimics the operation of a standard 100 Mbit ethernet bus interface so that the MAC functionality can be exploited. The programmable logic device contains a state machine that performs bus arbitration.

10 Figure 17 shows the state diagram for the state machine in the programmable logic device for the management bus. The arbitration lines determine which module has control of the bus by using open collector logic. The pull-ups for the arbitration bus reside on the chassis controller modules. Each 15 slot places its slot ID on the arbitration lines to request the bus. During transmission of the preamble of data to be transmitted, if the arbitration is corrupted, the bus controller assumes that another slot had concurrently requested the bus and the state machine within the PLD aborts 20 the transfer operation by forcing a collision signal for both the bus and the local MAC device active. As other modules detect the collision signal on the bus active, the collision line on each local MAC is forced to the collision state, which allows the back-off algorithm within the MAC to determine the 25 next transmission time. If a collision is not detected, the data is latched into the receive FIFO of each module, and the TX_Enable signal is used to quantify data from the bus. The state machine waits four clock cycles during the preamble of the transmit state, and four clock cycles during the collision 30 state to allow the other modules to synchronize to the state of the bus.

Backplane Architecture

Fig. 7 shows the internal backplane architecture of the current embodiment of the switch of the invention that was 35 shown in exemplary fashion in Fig. 6. One feature is the full-mesh interconnection between slots shown in the region 505. Slots are shown by vertical lines in Figure 7. This is implemented using 144 pairs of differential pairs embedded in

the backplane as shown in Figure 8. Each slot thus has a full-duplex serial path to every other slot in the system. There are $n(n-1)$ non-looped-back links in the system, that is, 132 links, doubled for the duplex pair configuration for a 5 total of 264 differential pairs (or further doubled for 528 wires) in the backplane to create the backplane mesh in the present embodiment of the invention. Each differential pair is able to support data throughput of more than 1 gigabit per second.

10 In the implementation of the current invention, the clock signal is embedded in the serial signaling, obviating the need for separate pairs (quads) for clock distribution. Because the data paths are independent, different pairs of cards in the chassis may be switching (ATM) cells and others switching 15 (IP) packets. Also, each slot is capable of transmitting on all 11 of its serial links at once, a feature useful for broadcasting. All the slots transmitting on all their serial lines achieve a peak bandwidth of 132 gigabits per second. Sustained bandwidth depends on system configuration.

20 The mesh provides a fully redundant connection between the application cards in the backplane. One connection can fail without affecting the ability of the cards to communicate. Routing tables are stored in the chassis controllers. If, for example, the connection between 25 application module 1 and application module 2 failed, the routing tables are updated. The routing tables are updated when the application modules report to the chassis controllers that no data is being received on a particular serial link. Data addressed to application module 2 coming in through 30 application module 1 is routed to another application module, for instance application module 3, which would then forward the data to application module 2.

The bus-connected backplane region 525 includes three bus systems. The management/control bus 530 is provided for out-35 of-band communication of signaling, control, and management information. A redundant backup for a management bus failure will be the mesh interconnect fabric 505. In the current implementation, the management bus provides 32-bit 10-20 MHz

transfers, operating as a packet bus. Arbitration is centralized on the system clock module 102 (clock A). Any slot to any slot communication is allowed, with broadcast and multicast also supported. The bus drivers are integrated on

5 the System Bus FPGA/ASIC.

A TDM (Time Division Multiplexing) Fabric 535 is also provided for telephony applications. Alternative approaches include the use of a DSO fabric, using 32 TDM highways (sixteen full-duplex, 2048 FDX timeslots, or approximately 3

10 T3s) using the H.110 standard, or a SONET ATM (Asynchronous Transfer Mode) fabric.

Miscellaneous static signals may also be distributed in the bus-connected backplane region 540. Slot ID, clock failure and management bus arbitration failure may be

15 signaled.

A star interconnect region 545 provides independent clock distribution from redundant clocks 102 and 103. The static signals on backplane bus 540 tell the system modules which system clock and bus arbitration slot is active. Two clock

20 distribution networks are supported: a reference clock from which other clocks are synthesized, and a TDM bus clock, depending on the TDM bus architecture chosen. Both clocks are synchronized to an internal Stratum 3/4 oscillator or an externally provided BITS (Building Integrated Timing Supply).

Fig. 8 shows a first connection point on a first MCC on a first module, MCC A 350, and a second connection point on a second MCC on a second module, MCC B 352, and connections 354, 355, 356, 357 between them. The connections run through a backplane mesh 360 according to the present invention. There

25 are transmit 362, 364 and receive 366, 368 channels at each MCC 350, 352, and each channel has a positive and a negative connection. In all, each point on a module has four

30 connections between it and every other point due to the backplane mesh. The differential transmission line impedance and length are controlled to ensure signal integrity and high

35 speed operation.

Fig. 9 is a block diagram of the MCC chip. An F-bus interface 805 connects the MCC 300 to the FIFO bus (F-bus).

Twelve transmit FIFOs 810 and twelve receive FIFOs 815 are connected to the F-bus interface 805. Each transmit FIFO has a data compressor (12 data compressors in all, 820), and each receive FIFO has a data expander (12 data expanders in all, 825).
5 Twelve serializer/deserializers 830 serve the data compressors 820 and data expanders 825, one compressor and one expander for each. A channel in the MCC is defined as a serial link together with its encoding/decoding logic, transmit queue and receive queue. The serial lines running
10 from the channels connect to the backplane mesh. All the channels can transmit data at the same time. A current implementation of the invention uses a Mesh Communication Chip to interconnect up to thirteen F-buses in a full mesh using serial link technology. Each MCC has two F-bus interfaces and
15 twelve serial link interfaces. The MCC transmits and receives packets on the F-buses in programmable size increments from 64 bytes to entire packets. It contains twelve virtual transmit processors (VTPs) which take packets from the F-bus and send them out the serial links, allowing twelve outgoing packets
20 simultaneously. The VTPs read the MCC tag on the front of the packet and dynamically bind themselves to the destination
25 slot(s) indicated in the header.

The card/slot-specific processor, card/slot-specific MAC/PHY pair (Ethernet, SONET, HFC, etc.) and an MCC communicate on a bi-directional F-bus (or multiple unidirectional F-busses). The packet transmit path is from the PHY/MAC to the processor, then from the processor to the MCC and out the mesh. The processor does Layer 3 and Layer 4 look-ups in the FIPP to determine the packet's destination and
30 Quality of Service (QoS), modifies the header as necessary, and prepends the MCC tag to the packet before sending it to the MCC.

The packet receive path is from the mesh to the MCC and on to the processor, then from the processor to the MAC/Phy and out the channel. The processor strips off the MCC tag
35 before sending the packet on to the MAC.

A first data flow control mechanism in the present invention takes advantage of the duplex pair configuration of

the connections in the backplane and connections to the modules. The MCCs have a predetermined fullness threshold for the FIFOs. If a receive FIFO fills to the predetermined threshold, a code is transmitted over the transmit channel of the duplex pair to stop sending data. The codes are designed to direct-couple balance the signals on the transmission lines and to enable the detection of errors. The codes in the present implementation of the invention are 16B/20B codes, however other codes may be used within the scope of the present invention. The MCC sends an I1 or I2 code with the XOFF bit set to turn off the data flow. This message is included in the data stream transmitted on the transmit channel. If the FIFO falls below the predetermined threshold, the MCC clears the stop message by sending an I1 or I2 code with the XOFF bit cleared. The efficient flow control prevents low depth FIFOs from overrunning, thereby allowing small FIFOs in ASICs, for example, 512 bytes, to be used. This reduces microchip costs in the system.

Fig. 10 shows a packet tag, also called the MCC tag. The MCC tag is a 32-bit tag used to route a packet through the backplane mesh. The tag is added to the front of the packet by the slot processor before sending it to the MCC. The tag has four fields: a destination mask field, a priority field, a keep field, and a reserved field. The destination mask field is the field holding the mask of slots in the current chassis to which the packet is destined, which may or may not be the final destination in the system. For a transmit packet, the MCC uses the destination mask to determine which transmit queue(s) the packet is destined for. For a receive packet the MCC uses the priority and keep fields to determine which packets to discard in an over-committed slot. The reserved field is unused in the current embodiment of the invention.

The MCC has two independent transmit mode selectors, slot-to-channel mapping and virtual transmit mode. In slot-to-channel mapping, the MCC transparently maps SIDs to CIDs and software does not have to keep track of the mapping. In virtual transmit mode, the MCC handles multicast packets semi-transparently. The MCC takes a single F-bus stream and

directs it to multiple channels. The transmit ports in the MCC address virtual transmit processors (VTPs) rather than slots. The F-bus interface directs the packet to the selected virtual transmit processor. The VTP saves the Destination
5 Mask field from the MCC tag and forwards the packet data (including the MCC tag) to the set of transmit queues indicated in the Destination Mask. All subsequent 64-byte "chunks" of the packet are sent by the slot processor using the same port ID, and so are directed to the same VTP. The
10 VTP forwards chunks of the packet to the set of transmit queues indicated in the Destination Mask field saved from the MCC tag. When a chunk arrives with the EOP bit set, the VTP clears its destination mask. If the next chunk addressed to that port is not the start of a new packet (i.e., with the SOP
15 bit set), the VTP does not forward the chunk to any queue.

The destination mask of the MCC tag enable efficient multicast transmission of packets through "latching." The destination mask includes code for all designated destination slots. So, if a packet is meant for all twelve slots, only
20 one packet need be sent. The tag is delivered to all destinations encoded in the mask. If only a fraction of the slots are to receive the packet, only those slots are encoded into the destination mask.

The MCC maintains a set of "channel busy" bits which it
25 uses to prevent multiple VTPs from sending packets to the same CID simultaneously. This conflict prevention mechanism is not intended to assist the slot processor in management of busy channels, but rather to prevent complete corruption of packets in the event that the slot processor accidentally sends two
30 packets to the same slot simultaneously. When the VTPs get a new packet, they compare the destination CID mask with the channel busy bits. If any channel is busy, it is removed from the destination mask and an error is recorded for that CID. The VTP then sets all the busy bits for all remaining
35 destination channels and transmits the packet. When the VTP sees EOP on the F-bus for the packet, it clears the channel busy bits for its destination CIDs.

The F-bus interface performs the I/O functions between the MCC and the remaining portion of the application module. The application module adds a 32-bit packet tag (MCC tag), shown in Figure 10, to each data packet to be routed through 5 the mesh.

The data received or transmitted on the F-bus is up to 64 bits wide. In data transmission, the F-bus interface adds 4 status bits to the transmit data to make a 68-bit data segment. The F-bus interface drops the 68-bit data segment 10 into the appropriate transmit FIFO as determined from the packet tag. The data from a transmit FIFO is transferred to the associated data compressor where the 68-bit data segment is reduced to 10-bit segments. The data is then passed to the associated serializer where the data is further reduced to a 15 serial stream. The serial stream is sent out the serial link to the backplane.

Data arriving from the backplane comes through a serial link to the associated channel. The serializer for that channel expands the data to a 10-bit data segment and the 20 associated data expander expands the data to a 68-bit data segment which is passed on to the related FIFO and then from the FIFO to the F-bus interface.

A Fast IP Processor (FIPP) is provided with 32/64 Mbytes of high-speed synchronous SDRAM, 8 Mbytes of high-speed 25 synchronous SRAM, and boot flash. The FIPP has a 32-bit PCI bus and a 64-bit FIFO bus (F-bus). The FIPP transfers packet data to and from all F-bus-connected devices. It provides IP forwarding in both unicast and multicast mode. Routing tables are received over the management bus from the chassis route 30 server. The FIPP also provides higher layer functions such as filtering, and CoS/QoS.

Each line card has a clock subsystem that produces all the clocks necessary for each card. This will lock to the reference clock provided by the System Clock and Management 35 Bus Arbitration Card.

Each card has hot-plug, power-on reset circuitry, and Sanity Timer functions. All cards have on-board DC-to-DC converters to go from the -48V rail in the backplane to

whatever voltages are required for the application. Some cards (such as the CMTS card) likely will have two separate and isolated supplies to maximize the performance of the analog portions of the card.

5 Fig. 11 shows a generic switch header for the integrated switch. The header is used to route data packets through the system. The final destination may be either intra-chassis or inter-chassis.

10 The header type field indicates the header type used to route the packet through the network having one or more chassis systems. Generally, the header type field is used to decode the header and provide information needed for packet forwarding. Specifically, the header type field may be used to indicate that the Destination Fabric Interface Address has 15 logical ports. The header type field is also used to indicate whether the packet is to be broadcast or unicast. The header type field is used to indicate the relevant fields in the header.

20 The keep field indicates whether a packet can be dropped due to congestion.

The fragment field indicates packet fragmentation and whether the packet consists of two frames.

The priority field is used to indicate packet priority.

25 The encap type field is a one bit field that indicates whether further layer 2 processing is needed before the packet is forwarded. If the bit is set, L2 is present. If the bit is not set, L2 is not present.

30 The Mcast type field is a one bit field that indicates whether the packet is a broadcast or multicast packet. It may or may not be used depending on the circumstances.

The Dest FIA (Fabric Interface Address) type field indicates whether the destination FIA is in short form (i.e., <chassis/slot/port>) or in long form (i.e., <chassis/slot/port/logical port>). This field may or may not 35 be used depending on the circumstances. This field may be combined with the header type field.

The Src FIA type field is a one bit field that indicates whether the source FIA is in short form (i.e.,

<chassis/slot/port>) or in long form (i.e., <chassis/slot/port/logical port>). This field may or may not be used depending on the circumstances. This field may be combined with the header type field.

5 The data type field is an x-bit field used for application to application communication using the switch layer. The field identifies the packet destination.

10 The forwarding info field is an x-bit field that holds the Forwarding Table Revision is a forwarding information next hop field, a switch next hop, that identifies which port the packet is to go out, along with the forward_table_entry key/id.

15 The Dest FIA field is an x-bit field that indicates the final destination of the packet. It contains chassis/slot/port and sometimes logical port information. A chassis of value 0 (zero) indicates the chassis holding the Master Agent. A port value of 0 (zero) indicates the receiver of the packet is an application module. The logical port may be used to indicate which stack/entity in the card is to 20 receive the packet. All edge ports and ICL ports are therefore "1"-based.

The Src FIA field is an x-bit field that indicates the source of the packet. It is used by the route server to identify the source of incoming packets.

25 Figure 12 is a flow chart of the general packet forwarding process. When a packet is received at one of the application modules of the switch, the module examines the BAS header, if one is present, to determine if the packet was addressed for the chassis to which the module is attached. If 30 not, the application module looks up the destination chassis in the routing table and forwards the packet to the correct chassis. If the packet was addressed for the chassis, the application module examines the header to determine whether the packet was addressed to the module (or slot). If not, the 35 application module looks up the destination slot in the mapping table and forwards the packet to the correct application module. If the packet was addressed to the application module, the application module compares the

forwarding table ID in the header to the local forwarding table revision. If there is a match, the module uses the pointer in the header to forward the packet on to its next destination.

5 Unicast Traffic Received from an ICL Port

Figure 13 is a diagram of an incoming ICL packet. The packet has a BAS header, an encap field that may be either set or not (L2 or NULL), an IP field, and a data field for the data.

10 Figure 14 is a diagram of a header for the packet of Figure 12. The header type may be 1 or 2. A header type of 1 indicates an FIA field that is of the format chassis/slot/port both for destination and source. A header type of 2 indicates an FIA field of the format chassis/slot/port/logical port for 15 both destination and source. The keep field is not used. The priority field is not used. The fragment field is not used. The next hop filed is not used. The Encap field is zero or 1. The Mcast field is not used. The DST FIA type may be 0 or 1. The SRC FIA type may be zero or one. The BAS TTL field is not 20 used. The forward info field is used. And the DST and SRC FIA fields are used.

It is to be understood that the above-described embodiments are simply illustrative of the principles of the invention. Various and other modifications and changes may be 25 made by those skilled in the art which will embody the principles of the invention and fall within the spirit and scope thereof.

What is claimed is:

1. A process for switching data flow from one of a plurality of data communication application modules in a chassis interconnected across a backplane in said chassis, said process comprising the steps of:
 - A) receiving data at a one of said modules;
 - B) applying said data m bits at a time to an m -bit-wide parallel bus, where m is greater than one;
 - C) determining on which port of said one of said modules said data is to be transmitted, said port being directly connected through said backplane to another one of said modules to which said data is to be directed;
 - D) appending a tag to said data to indicate on which port said data is to be transmitted;
 - E) gating said m bits of data into a channel associated with said transmitting port; and
 - F) transmitting said data from said transmitting port across said backplane to said other one of said modules to which said data is to be directed.
2. The process of claim 1 further comprising the step of serializing said m bits of data before transmitting said data.
3. The process of claim 1 wherein said data is transmitted on a differential pair of connections.
4. The process of claim 1 wherein said transmission is self-clocked.
5. The process of claim 1 wherein said process proceeds simultaneously with and independently of the same steps being taken in another of said modules.
6. The process of claim 1 wherein it is determined in step (C) that said data is to be multicast through multiple ports of said one of said modules, each of said multiple ports being directly connected through said backplane to another one of said modules to which said data is to be directed; the tag in step (D) indicates multicasting according to a mask indicating multicast transmitting ports; and steps (E)-(F) occurs simultaneously for each of said multicast transmitting ports.
7. A meshed backplane in a chassis for interconnecting N data communication application modules for data communication

between them with N greater than 2, said meshed backplane comprising:

N(N-1) full-duplex serial links each connecting one of said N application modules with another with a minimum

5 throughput of 64 kilobits per second.

8. The backplane of Claim 7 wherein each of said full-duplex serial links comprises one differential pair dedicated to transmission from one of said application modules to another and one differential pair dedicated to receiving to said one 10 of said application modules.

9. The backplane of Claim 7 further comprising:

a management information bus providing a connection to each application module and to a chassis control module.

10. The backplane of Claim 7 further comprising:

15 at least two redundant management information buses each providing a connection to each application module and to at least two chassis control modules.

11. The backplane of Claim 7 further comprising a power distribution system.

20 12. The backplane of Claim 7 further comprising a telephony mesh.

13. A data switch, comprising:

a backplane mesh organized as a plurality of slots to connect to a plurality of application modules, said backplane 25 mesh providing an independent data path from each said slot to every other slot, said backplane mesh providing a loopback in each said slot;

a controller module having a routing table; and

30 a bus connected to each said slot and said controller module, said bus providing control data for the data flow through the backplane mesh,

whereby said independent data paths enable different types of data traffic to flow simultaneously through through said data switch.

35 14. The data switch of claim 13 wherein each said slot has a slot ID.

15. The data switch of claim 13 further comprising:

at least one mesh communication chip (MCC) having a plurality of channels,

 said slots having a plurality of ports, each said channel in said MCC linking to a port.

5 16. The data switch of claim 15 further comprising:
 a slot ID for each said slot; and,
 means for said MCC to transmit said slot ID,
 whereby flexible routing of said backplane is enabled.

10 17. The data switch of claim 15, wherein said MCC further comprises:
 a plurality of FIFOs, one FIFO for each channel;
 means for receiving a packet having a destination mask;
 means for reading said destination mask;
 means for copying said packet to appropriate ones of said
15 plurality of FIFOs in response to said destination mask,
 whereby said packet may be substantially simultaneously transmitted over a plurality of channels.

18. A method for flexibly routing a backplane, comprising the steps of:
20 providing a backplane mesh organized as a plurality of slots to connect to a plurality of application modules;
 providing each said slot with a slot ID;
 reading, upon insertion of each said plurality of application modules into a slot of said plurality, the slot ID
25 of said slot;
 transmitting from each application module the slot ID of the slot to which said application module is connected; and,
 building a routing table in response to each said transmitted slot IDs,
30 whereby said backplane may be flexibly routed.

19. The method of claim 18 further comprising the step of controlling differential line impedance and length
 whereby signal integrity is maintained and high speed operation is enabled.

35 20. The method of claim 18 further comprising the step of building a mapping table in each said application module.

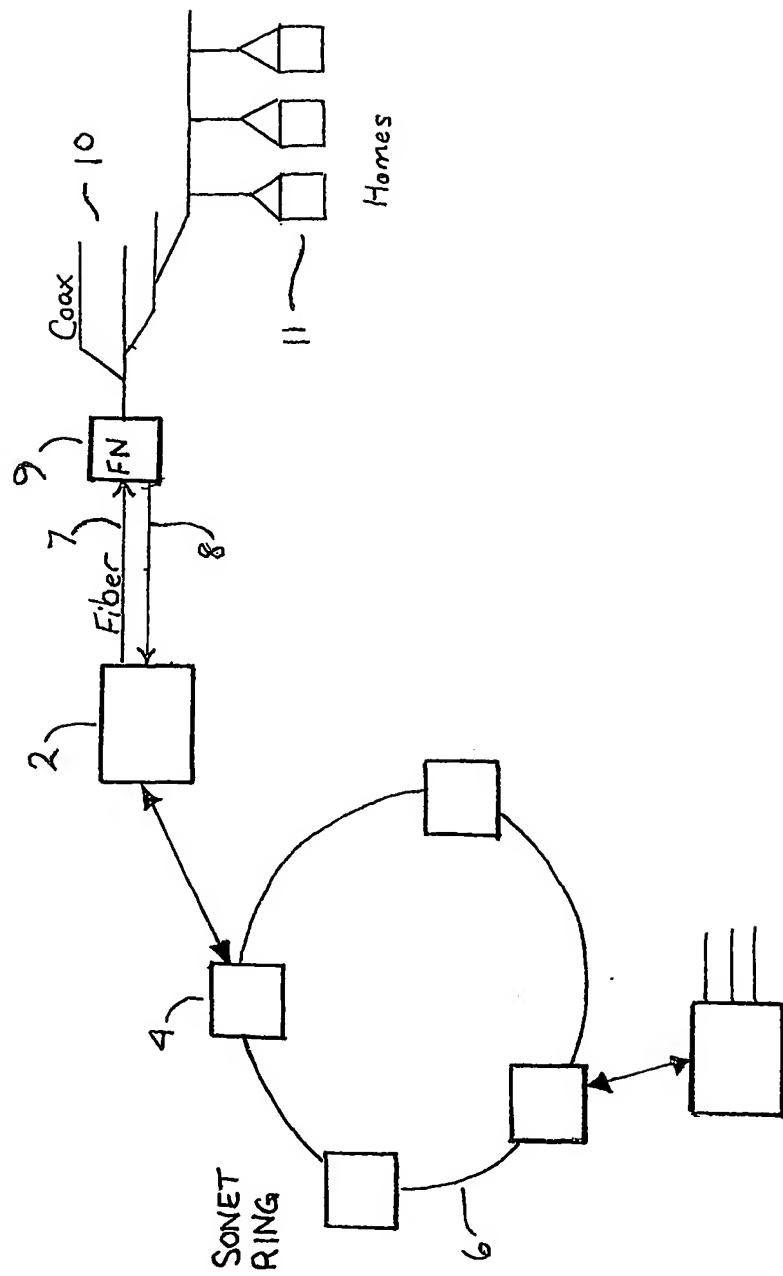


FIG. 1
Prior Art

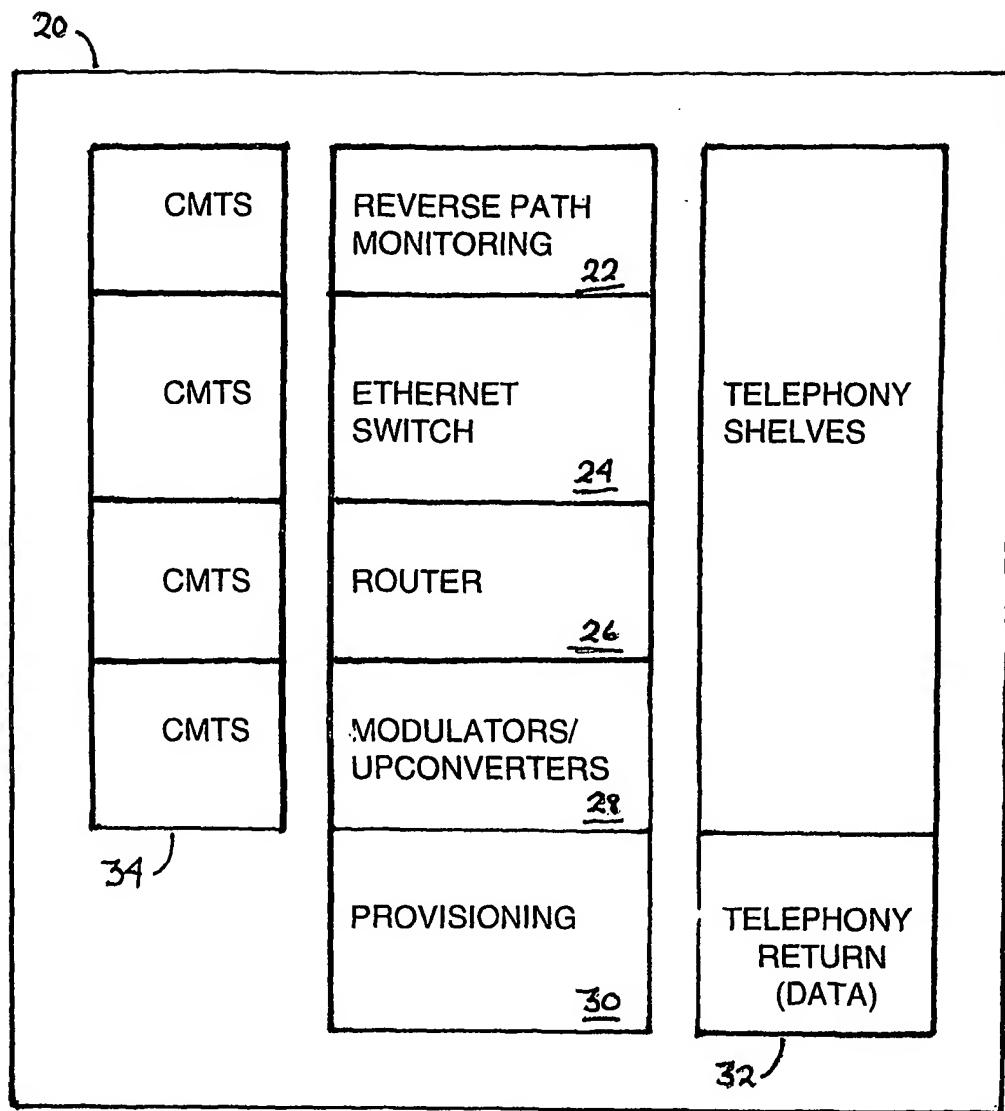


FIG. 2
(PRIOR ART)

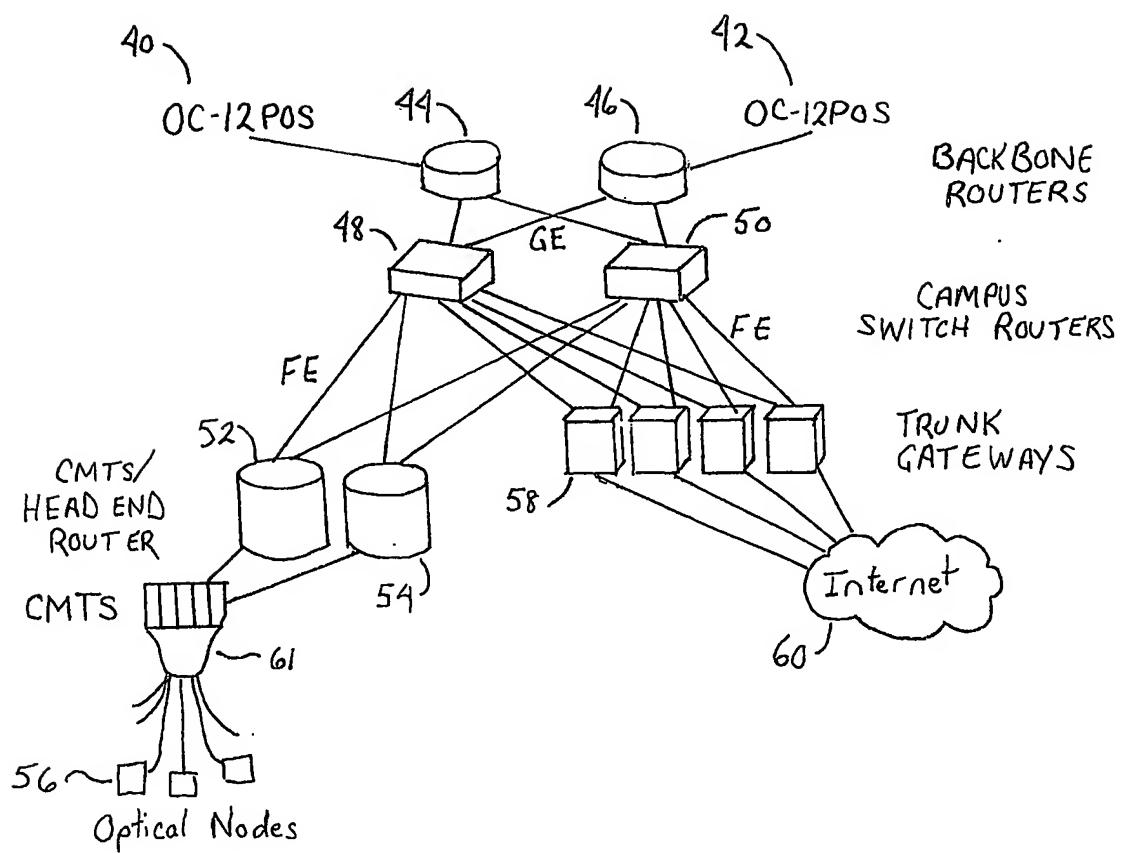


FIG. 3
Prior Art

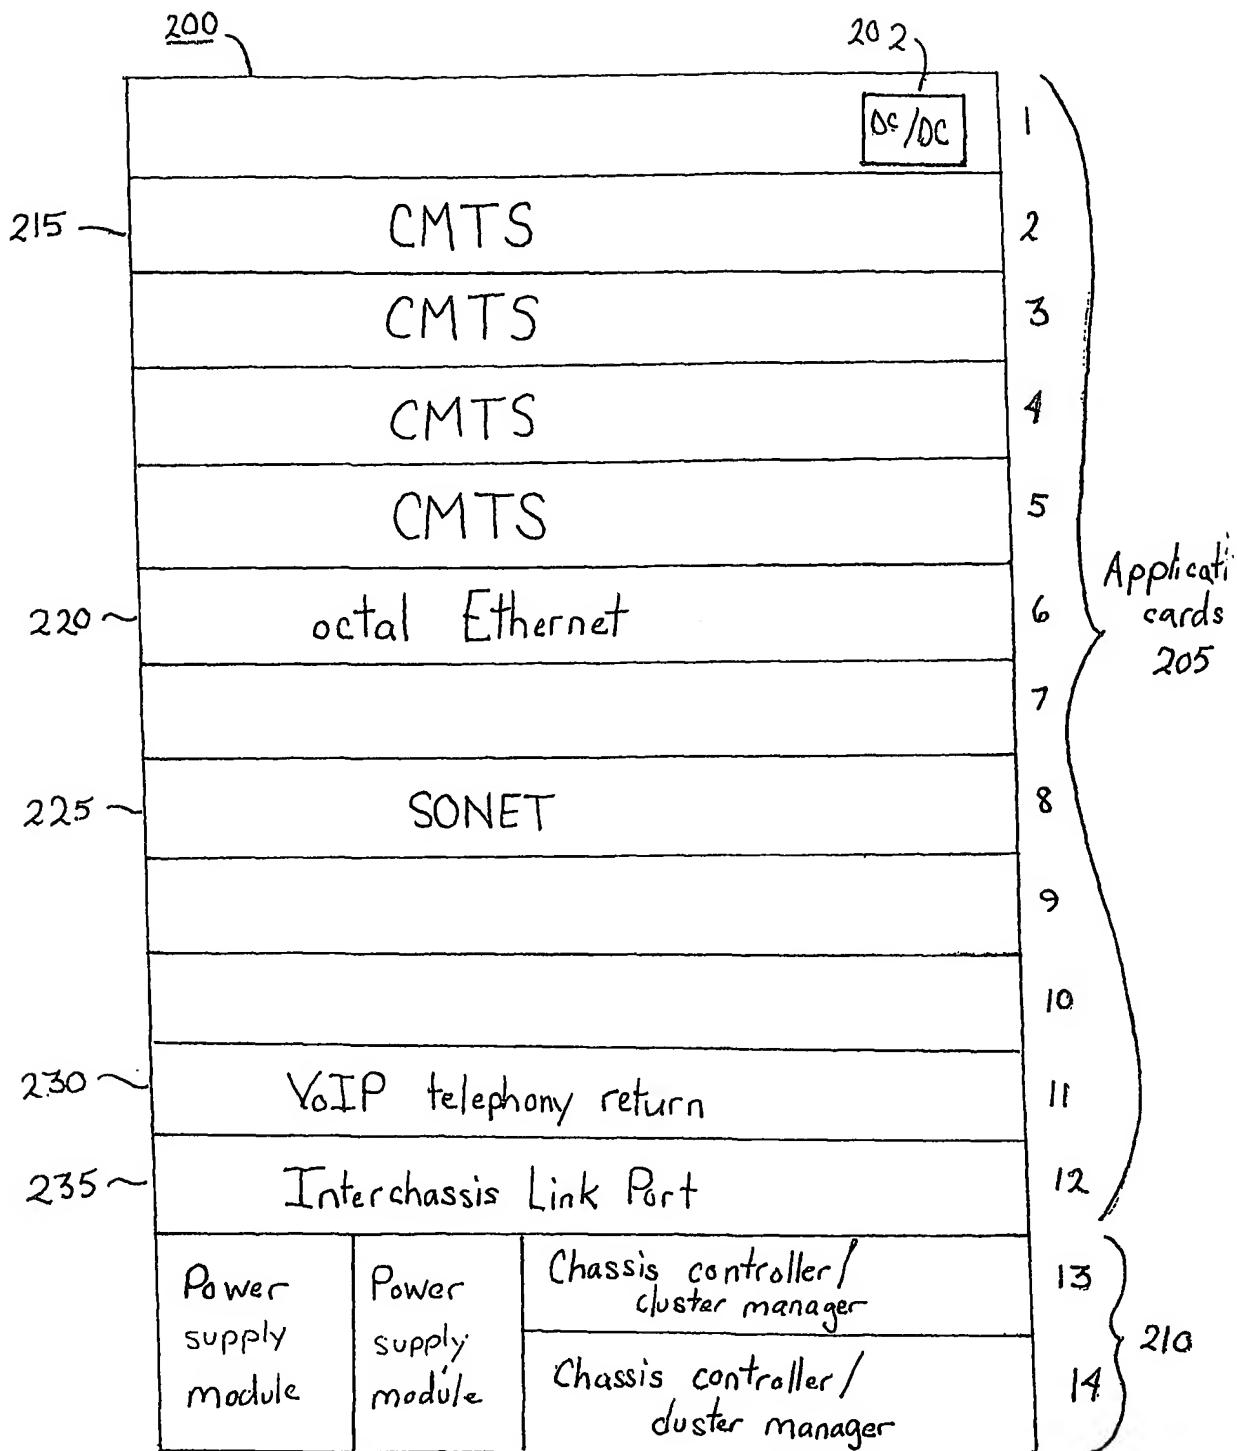


FIG. 4

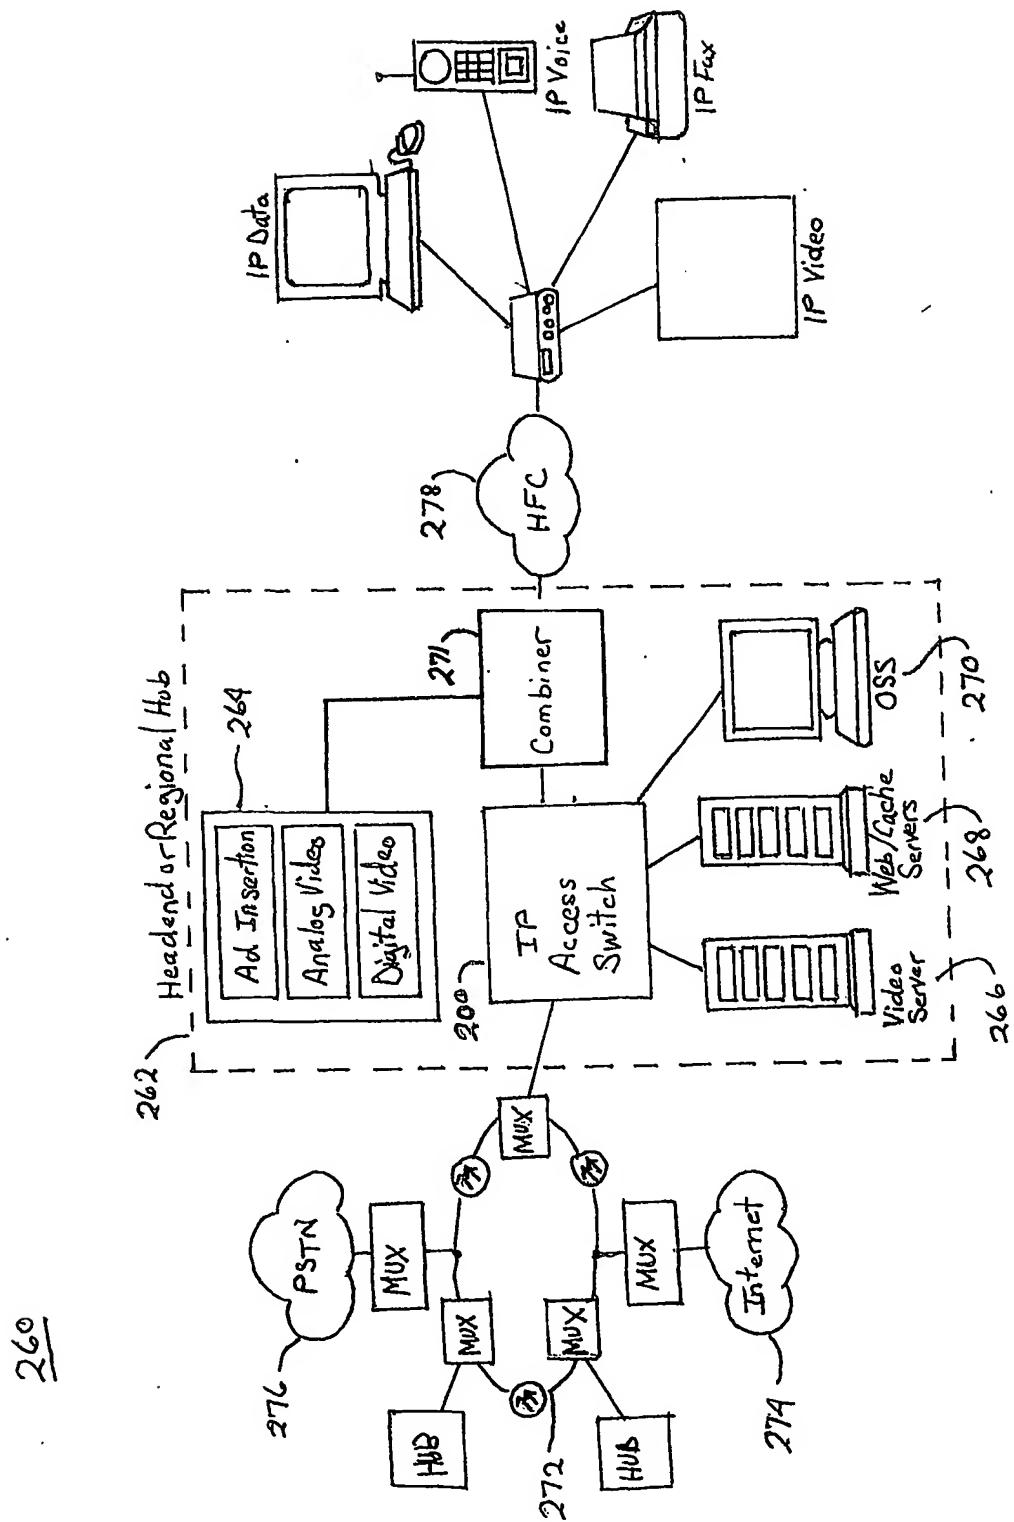


FIG. 5

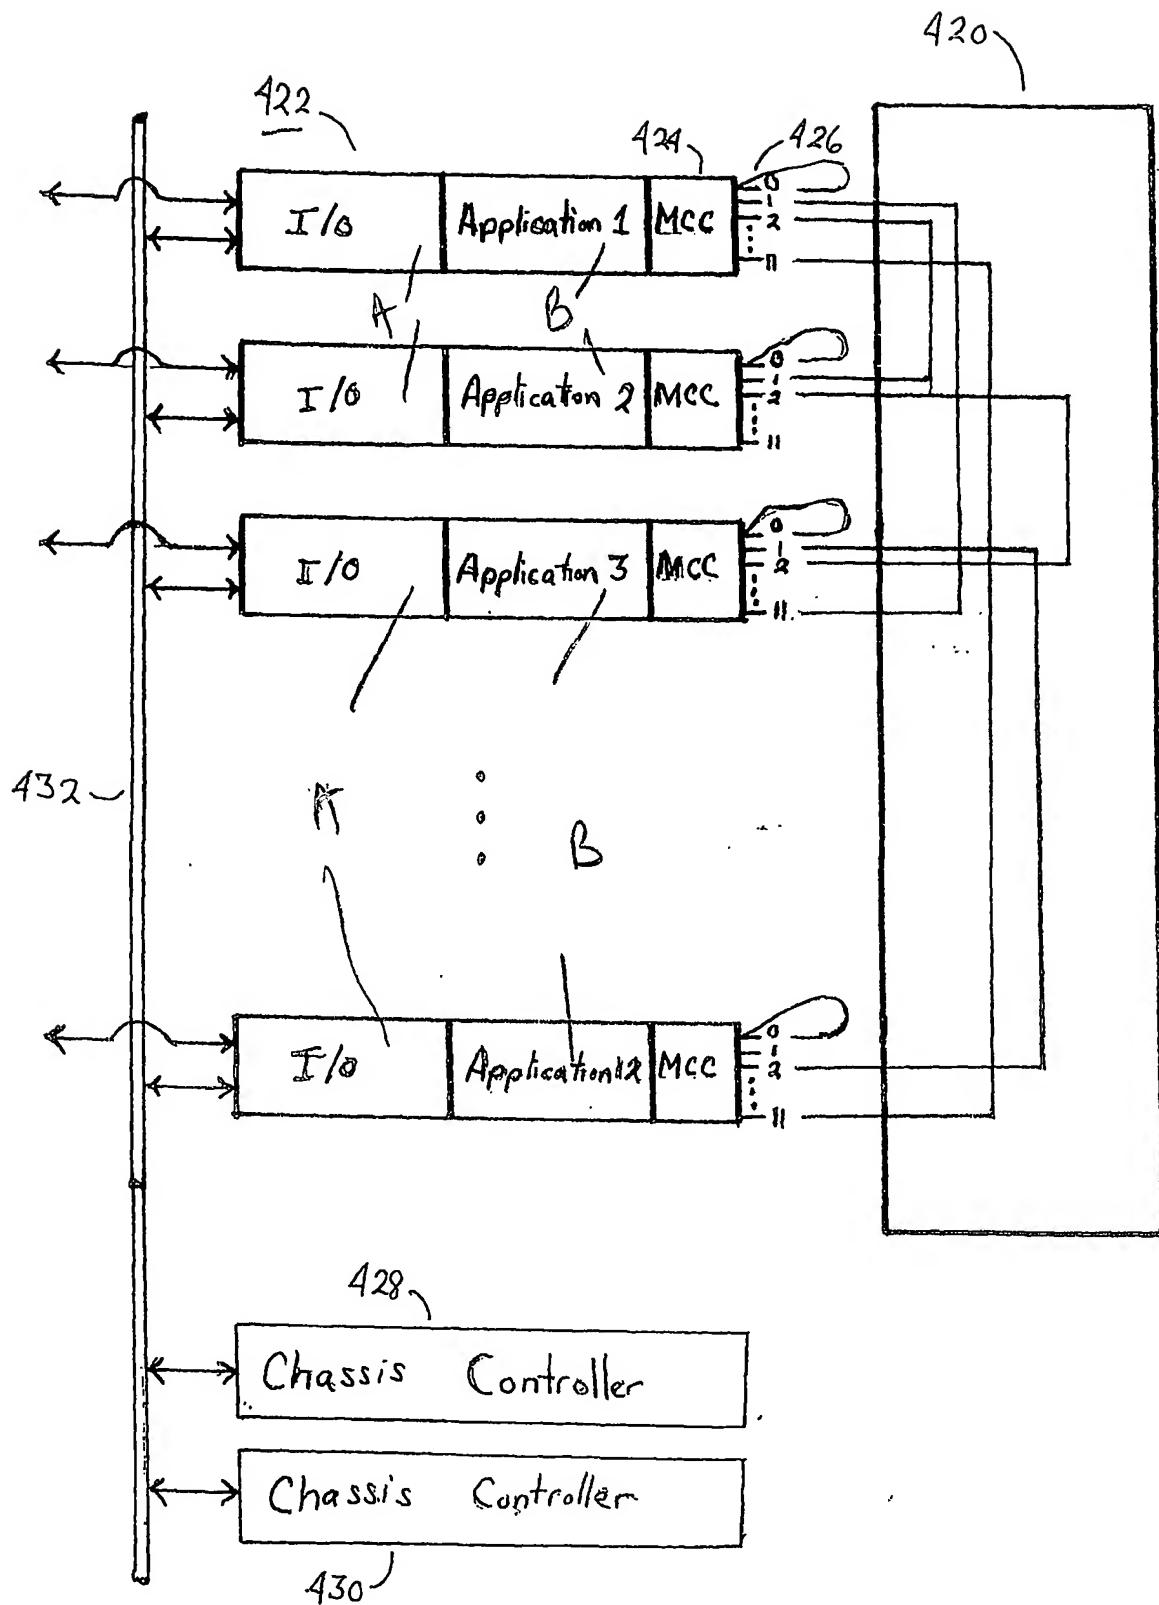


FIG. 6

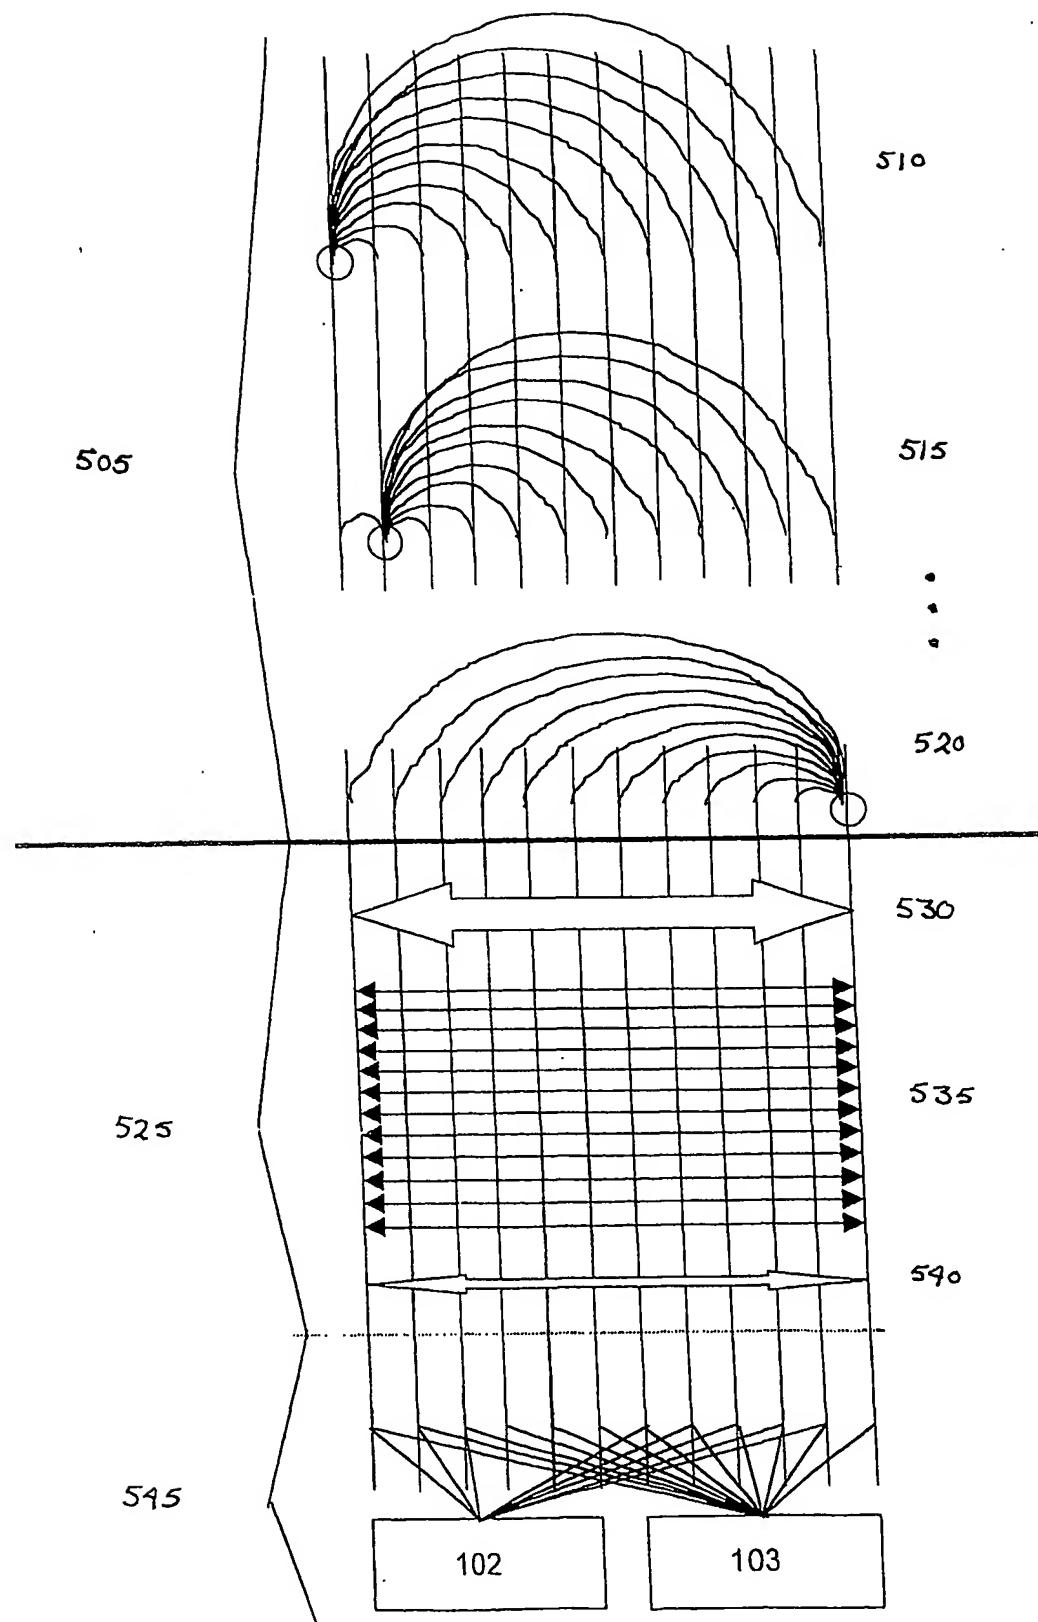


FIG. 7

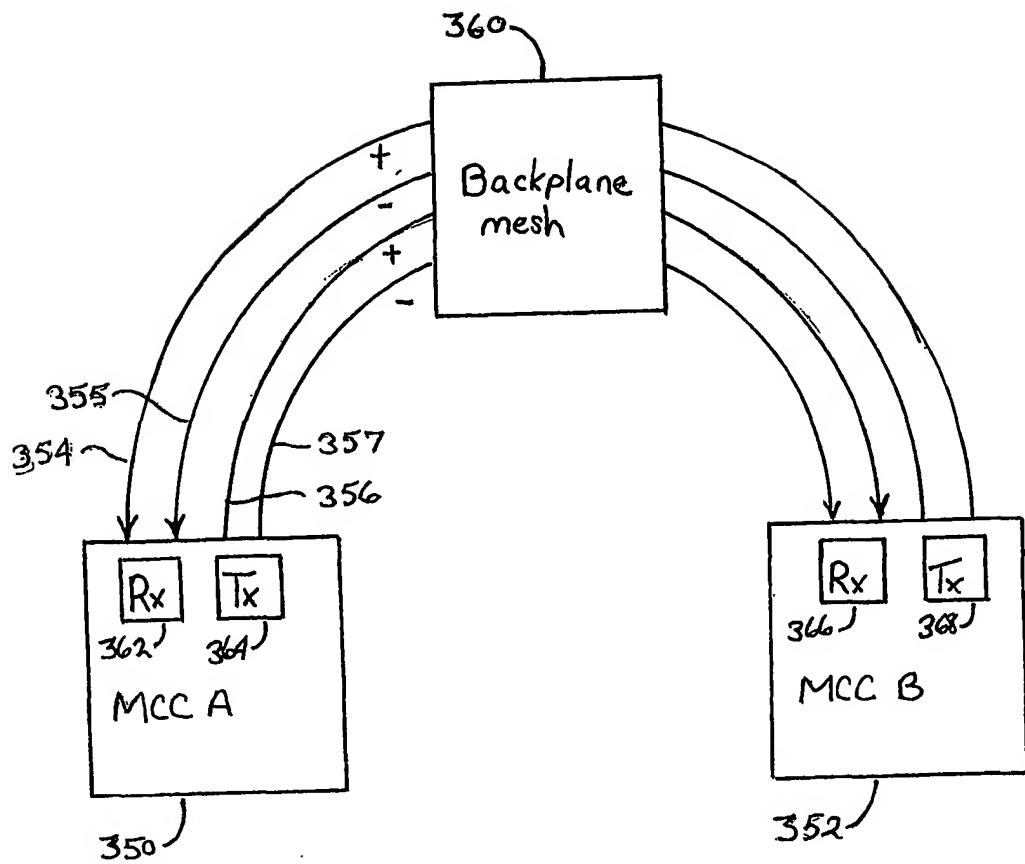


FIG. 8

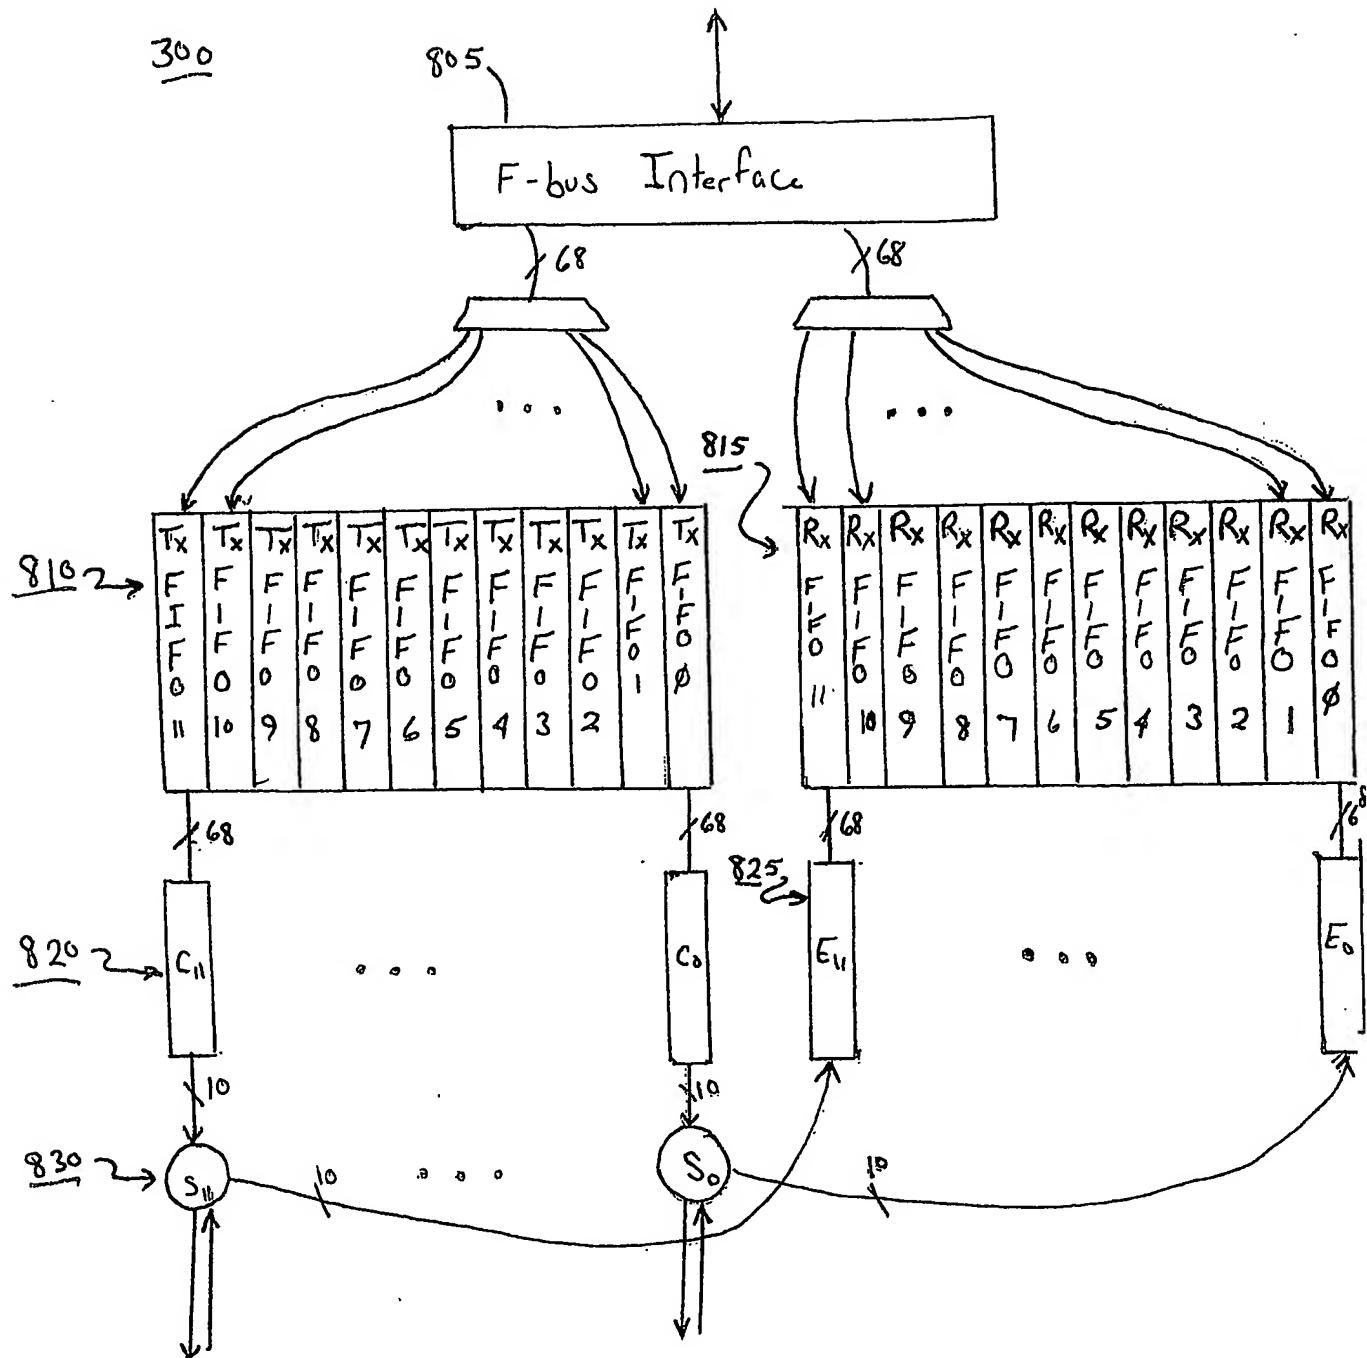


FIG. 9

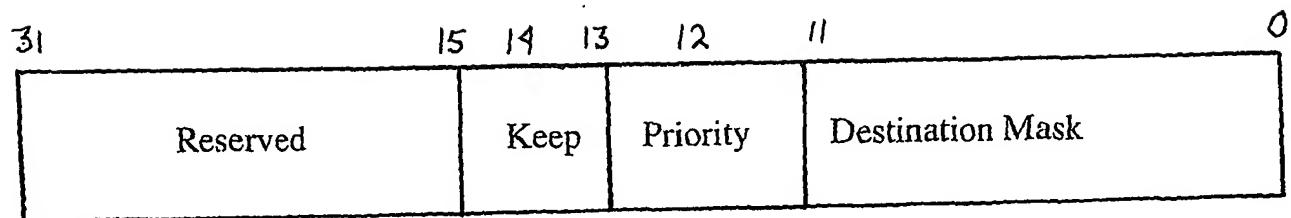
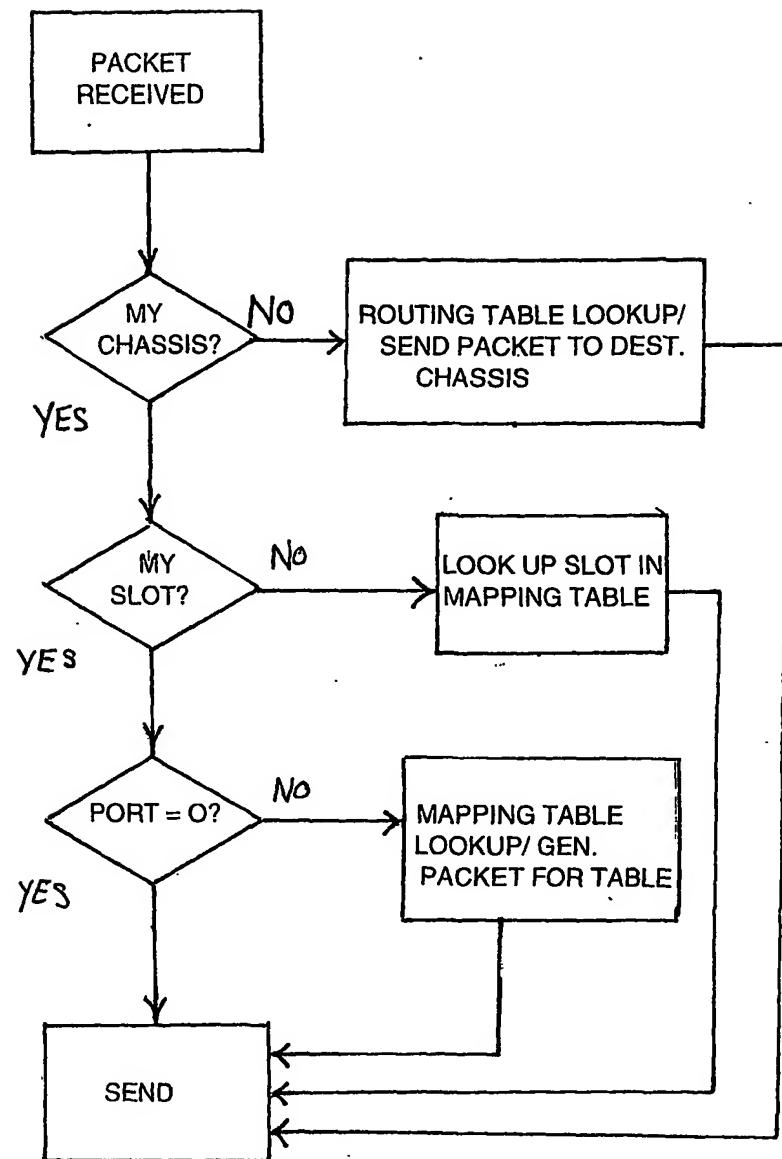


FIG. 10

Header Type	Keep	Fragment	Priority	Encap Type
Mcast Type	Dest FIA Type	Src FIA Type	Data Type	Forwarding Info
Dest FIA	Src FIA			

FIG. 11

**FIG. 12**

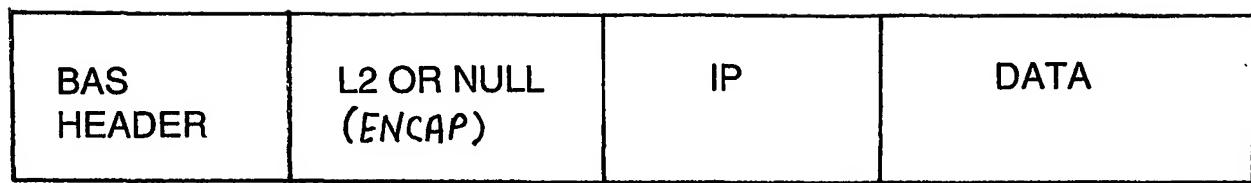


FIG. 13

BAS TYPE 1 OR 2	KEEP	PRIORITY	FRAGMENT	NEXT HOP	ENCAP	MCAST	DST FIA TYPE	SRC FIA TYPE
BAS TTL	FOR- WARD INFO		DST/SRC FIA					

FIG. 14

	Destination Slot											
	0	1	2	3	4	5	6	7	8	9	10	11
0	0	5	2	4	7	3	1	11	6	8	9	10
1	10	0	7	5	3	4	2	11	1	6	8	9
2	3	10	0	7	4	2	5	1	11	8	6	9
3	1	9	10	0	7	4	3	2	6	8	11	5
4	2	3	1	10	0	7	4	6	8	5	9	11
5	11	2	5	1	10	0	7	8	4	9	3	6
6	11	5	6	1	9	10	0	7	4	3	8	2
7	9	11	1	2	6	8	10	0	7	4	5	3
8	9	11	1	5	8	3	2	10	0	7	6	4
9	1	6	5	9	8	10	3	2	11	0	7	4
10	6	8	5	9	10	3	2	4	1	11	0	7
11	7	5	9	6	8	10	4	3	2	1	11	0

TRANSMIT CHANNELS (N SIDE)

FIG. 15

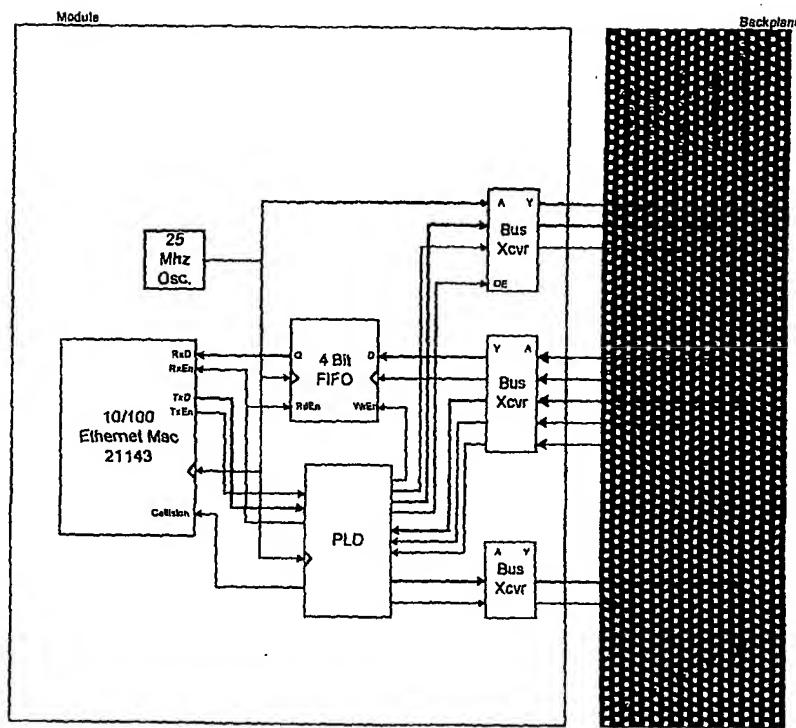


FIG. 16

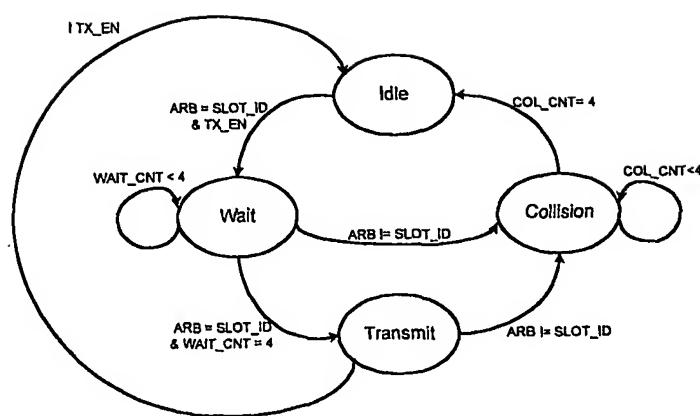


FIG. 17